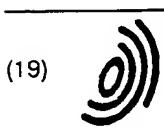


W204-0



(19)

Europäisches Patentamt

European Patent Office

Office européen des brevets



(11) EP 1 026 597 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
09.08.2000 Bulletin 2000/32(51) Int. Cl.⁷: G06F 13/40, G06F 13/28

(21) Application number: 00300446.2

(22) Date of filing: 21.01.2000

(84) Designated Contracting States:
 AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
 MC NL PT SE
 Designated Extension States:
 AL LT LV MK RO SI

(30) Priority: 03.02.1999 US 243169

(71) Applicant:
 SUN MICROSYSTEMS, INC.
 Palo Alto, California 94303 (US)

(72) Inventor: Crosland, Andrew
Haddenham, Aylesbury HP17 8EX (GB)

(74) Representative:
 Harris, Ian Richard et al
 D. Young & Co.,
 21 New Fetter Lane
 London EC4A 1DA (GB)

(54) Data alignment between buses

(57) A bus bridge for a computer system for bridging first and second buses includes a shift and accumulate unit. The shift and accumulate unit includes a shifter having an input connected to receive bytes from one of the first and second buses and an output providing a selectable shift to the received bytes. The shift and accumulate unit also includes an accumulator having an input connected to receive the output of the shifter and providing accumulation of selectable bits of the shifted bytes, the accumulator having an output for supplying realigned bytes to be passed to the other of the first and second buses. The combination of the shifter and the accumulator permits a desired amount of shift to be combined with the accumulation of selected bits or bytes to realign sets of bytes from one bus and to form sets of bytes for the other bus. Burst transfer is also possible by operating the shift and accumulate unit to operate in successive cycles for successive sets of input bytes from one of the buses.

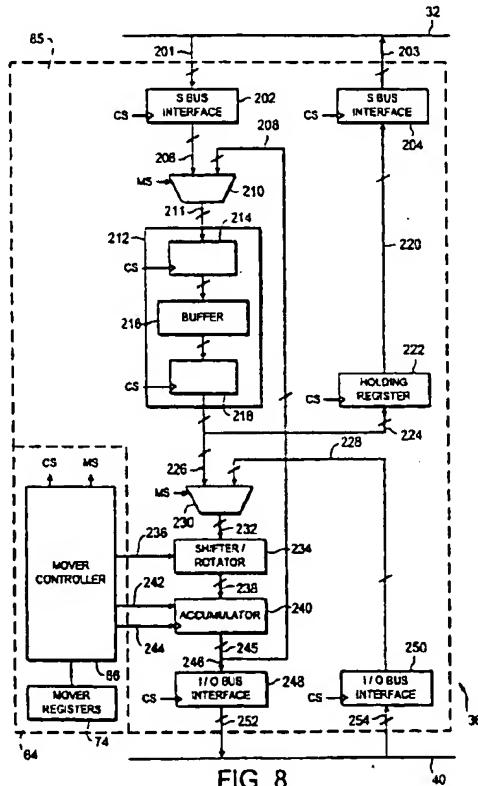


FIG. 8

EP 1 026 597 A2

BEST AVAILABLE COPY

Description**BACKGROUND OF THE INVENTION**

[0001] The invention relates to transferring data between buses in a computer system. In particular, the invention relates to a bus bridge for transferring data between buses, to a DMA controller for transferring data between buses, to a computer system incorporating such a bus bridge and/or DMA controller and to a method of transferring data between buses.

[0002] In the present context, a bridge is a mechanism for transferring data between buses. It is commonplace in modern computer systems for multiple buses to be provided, and for a bridge mechanism to be provided for transferring data between the buses. The bridge mechanism may take many different forms and may incorporate other functions. For example, it is known to provide a DMA controller between an I/O bus and another bus forming part of the computer system.

[0003] For different buses, it may be desirable to have different byte address alignments. The present invention addresses the problem of providing efficient transfer between buses with selectable bus alignment on the respective buses.

SUMMARY OF THE INVENTION

[0004] Particular and preferred aspects of the invention are set out in the accompanying independent and dependent claims. Combinations of features from the dependent claims may be combined with features of the independent claims as appropriate and not merely as explicitly set out in the claims.

[0005] In accordance with one aspect of the invention, there is provided a bus bridge for a computer system. The computer system can include a first bus and a second bus with the bridge being locatable between the first and second buses. The bridge includes a shifter having an input connected to receive bytes from one of the first and second buses and an output providing a selectable shift to the received bytes. It also includes an accumulator having an input connected to receive the output of the shifter and providing selective accumulation of the shifted bytes, the accumulator having an output for supplying realigned bytes to be passed to the other of the first and second buses.

[0006] The combination of the shifter and the accumulator permits a desired amount of shift to be combined with the accumulation of selected bytes, or bits, to realign sets of bytes from one bus and to form sets of bytes for the other bus. This structure enables flexible transfer of bytes with an arbitrary realignment in an efficient manner. Burst transfer is also possible by operating the shifter and the accumulator unit in successive cycles for successive sets of input bytes input from the input bus.

[0007] Control logic provides control signals to

select the shift applied to the received bytes and to select the bytes, or bits, to be accumulated from the shifted bytes. The shifter and the accumulator are operable in successive phases within a cycle. In one cycle, a first set of bytes, or bits, from the shifted bytes may be selected for one set of output bytes. In another cycle, a second set of bytes, or bits, may be selected from the shifted bytes for another set of output bytes.

[0008] The control logic is arranged to determine the control signals for successive cycles and phases to effect a burst transfer of bytes from one bus to the other bus.

[0009] The accumulator can include an output register. In the first phase, accumulation of an output set of realigned bytes may be completed by registering selected parts of a shifted input set of bytes at corresponding parts in the output register. In, the second phase, accumulation of a subsequent output set of realigned bytes may be started by registering remaining parts of a shifted input set of bytes at corresponding parts in the output register.

[0010] It should be noted that a set of bytes may form a word, or may form part of a word, depending on the relative bus widths and the width of the shift and accumulate unit.

[0011] In an embodiment of the invention, the bridge further comprises a buffer for buffering sets of bytes, the buffer being connectable between the first bus and the shifter. The buffer can be connected between the first bus and the shifter when bytes are to be transferred from the first bus to the second bus, or between the accumulator and the first bus when bytes are to be transferred from the second bus to the first bus. This provides for flexible buffering of data to take account of possible different speeds of operation of the timings for the first and second buses.

[0012] In order to provide the selected ordering of the buffer unit and the shift and accumulate unit as described above, the buffer and the shifter are preferably associated with input multiplexors.

[0013] The input multiplexor of the buffer unit can include a first input connectable to receive bytes from the first bus, a second input connected to receive bytes from the accumulator and an output connected to output bytes to a buffer of the buffer unit.

[0014] The input multiplexor of the shifter can comprise a first input connected to receive bytes from the second bus, a second input connected to receive bytes from an output of the buffer unit and an output connected to output bytes to the shifter.

[0015] An embodiment of the bus bridge forms a direct memory access controller, for example for connection between an I/O bus and another processor bus, for example a bus operable under an appropriate system bus (SBus) protocol.

[0016] According to another aspect of the invention, therefore, there is provided a direct memory access controller for a computer system having a first bus and

BEST AVAILABLE COPY

a second bus. The direct memory access controller is locatable between the first and second buses. It comprises a shifter having an input connected to receive bytes from one of the first and second buses and an output providing a selectable shift to the received bytes. It also comprises an accumulator having an input connected to receive the output of the shifter and providing selective accumulation of the shifted bytes, the accumulator having an output for supplying realigned bytes to be passed to the other of the first and second buses.

[0017] In accordance with a further aspect of the invention, there is provided a computer system having a first bus, a second bus and a bus bridge as set out above between the first and second buses.

[0018] In accordance with yet a further aspect of the invention, there is provided a method of transferring data between a first bus and a second bus in a computer system. The method comprises steps of:

shifting bytes received from one of the first and second buses to providing a selectable shift to the received bytes; and

performing selective accumulation of the shifted bytes for supplying realigned bytes to be passed to the other of the first and second buses.

[0019] The steps are repeated in successive cycles, with a set of bytes being received in each cycle. The set of bytes can be a word, or part of a word as described above.

[0020] In a first phase of a cycle, accumulation of an output set of realigned bytes can be completed by registering bytes, or bits, from selected positions of a shifted input set of bytes at corresponding positions in an output register; and

in a second phase of a cycle, accumulation of a subsequent output set of realigned bytes can be started by registering bytes, or bits, from remaining positions of a shifted input set of bytes at corresponding positions in the output register

[0021] For transfers from the first bus to the second bus, bytes are buffered prior to being processed by the shifter and the accumulator, and, for transfers from the second bus to the first bus, bytes are buffered after being processed by the shifter and the accumulator.

[0022] For transferring a burst of bytes between the first and second buses, the method can include one or more of the following steps:

in an initial stage, selected bits of shifted bytes from up to two cycles are accumulated to form an initial set of output bytes;

in an intermediate stage, selected bits of shifted bytes from two cycles are accumulated to form intermediate sets of output bytes; and

in a final stage, selected bits of shifted bytes from

up to two cycles are accumulated to form a final set of output bytes.

[0023] For transferring a burst of bytes between the first and second buses, wherein, to achieve realignment with efficient block transfer, a block transfer may include one or more of the following stages:

- an initial transfer of selected bits;
- an initial transfer of selected bytes;
- an initial transfer of selected blocks of bytes;
- an initial transfer of sub-blocks of words;
- a block transfer of words;
- a final transfer of sub-blocks blocks of words;
- a final transfer of selected blocks of bytes;
- a final transfer of selected bytes; and
- a final transfer of selected bits.

BRIEF DESCRIPTION OF THE DRAWINGS

[0024] Exemplary embodiments of the present invention will be described hereinafter, by way of example only, with reference to the accompanying drawings in which like reference signs relate to like elements and in which:

Figure 1 is a schematic block diagram of an exemplary configuration of a computer system incorporating the present invention;

Figure 2 is a schematic block diagram of a DMA controller of the system of Figure 1;

Figure 3 is a schematic representation of a command buffer in main memory;

Figure 4 is a representation of a base register;

Figure 5 is a representation of a tail register;

Figure 6 is a representation of a head register;

Figures 7A and 7B are flow diagrams illustrating the operation of the computer system of Figure 1 for transferring DMA commands to the DMA controller;

Figure 8 is a schematic block diagram of a DMA transfer mechanism of the DMA controller Figure 2;

Figure 9 is a schematic block representation of a shifter/rotator of the mechanism of Figure 8;

Figure 10 is a schematic representation of an accumulator of the mechanism of Figure 8;

Figure 11 illustrates source and destination byte alignments for a data transfer;

Figure 12 illustrates operations of the shifter/rotator and accumulator;

Figure 13 is a flow diagram illustrating control operations for data transfer;

Figure 14 illustrates possible block transfer operations;

Figure 15 illustrates the configuration of the mechanism of Figure 8 for transfers between a first bus and a second bus; and

Figure 16 illustrates the configuration of the mechanism of Figure 3 for transfers between the second

bus and the first bus.

DESCRIPTION OF PARTICULAR EMBODIMENTS

[0025] An embodiment of the present invention is described hereinafter with reference to an example configuration of a computer system that is suitable for forming a part of a fault-tolerant computer system. However, it should be noted that the invention is not limited to such an embodiment, and that it may be implemented in any computer system employing a DMA controller.

[0026] Figure 1 is a schematic overview of an example of a computer system 10. The computer system 10 shown in Figure 1 includes first and second processors 12 and 14, in the present example CPU 1 and CPU 2. In an alternative embodiment, there may, for example, only be a single CPU. The processors 12 and 14 are connected to a first bus 22 operable under a UPA (Universal Processor Architecture) protocol. Also connected to the first bus 22 via a memory interface 20 is a main memory 16. A system I/O controller 30 acts as a bridge between the first bus 22 and a system bus (SBus) 32. A system controller 18 is connected to the first bus 22, and also, via control lines 24, 26, 28 and 29 to the processor 12, the processor 14, the main memory 16, and the system I/O controller 30, respectively.

[0027] One or more slots (e.g., slot 34) could be provided for the connection of further components (not shown) to the SBus 32. A processor I/O bus bridge 36 is connected to the SBus 32. This provides a connection to a peripheral bus (PBus) 33 for the connection of one or more peripherals 50. The processor I/O bridge 36 also provides for processor I/O access to one or more peripheral devices 46 and 48 via an I/O bus 40.

[0028] A direct memory access controller (DMAC) 38 is also connected between the I/O bus 40 and the SBus 32 to provide DMA access for the peripheral devices on the bus 40 to the processor 10, in particular to the main memory 16 of the processor 10. The DMAC 38 forms a further bus bridge.

[0029] For a fault tolerant computer system, the I/O bus could be configured as a pair of fault tolerant redundant buses, with multiple instances of the computer system 10 being connected thereto.

[0030] Figure 2 illustrates the DMAC 38 in more detail. The DMAC 38 comprises a loader machine 62 (hereinafter loader 62) and a mover machine 64 (hereinafter mover 64). Each of the loader 62 and the mover 64 has a processor I/O interface (66 and 68, respectively) which enables the processor 10 to control the loader and mover directly.

[0031] In an embodiment of the invention, the loader 62 comprises a loader controller 84 which enables the loader 62 to load DMA commands from the main memory 16 into the DMAC 38. In particular, the loader controller 84 is operable to load DMA commands from the main memory 16 of the processor 10 into mover registers 74 in the mover 64.

[0032] The loader 62 also comprises registers 72. These include, in particular, a base register 76 for a base pointer 75 to address a command buffer 92 in main memory 16, a tail pointer register 78 for a tail pointer 77 to address the tail of a sequence of DMA commands in the command buffer 92, and a head pointer register 80 for a head pointer 79 to address a head of a sequence of commands in the command buffer 92. Further registers 82 are provided in the loader 62.

[0033] The mover 64 is operable to move bulk data between the buses. To this effect, it includes a mover transfer mechanism 85. The mover 64 comprises a mover controller 86, which is operable to control DMA operations according to DMA commands in the mover registers 74. A DMA command for the mover 64 is a 16-byte structure set up by the processor 10. A DMA command causes the mover 64 to move data in either direction between the I/O bus 40 and an SBus DVMA (Direct Virtual Memory Access) address, which translates to main memory. The DMA command is either written directly into the mover registers 74 by the processor 10, or can be written to the command buffer in the main memory 16 and then be loaded into the mover registers 74 by the loader 62.

[0034] A DMA command can specify any start address for DVMA and fault tolerant I/O bus accesses. It can specify any transfer length up to a given maximum, with no restrictions on alignment. A DMA command can include:

- a DVMA address for main memory;
- an I/O bus address;
- a count of bytes; and
- a DVMA action byte.

[0035] The DVMA action byte comprises command and status information including a direction indicator (for the direction of transfer), a validity bit, an interrupt enable bit and a writeback enable bit.

[0036] In order to minimize the overhead for transferring DMA commands between the processor 10 and the DMAC 38, an embodiment of the present invention includes a particular mechanism to achieve the efficient transfer of DMA commands.

[0037] As illustrated in Figure 3, an area of main memory is configured as a command buffer 92. This area of main memory 16 configured as the command buffer 92 is located in the main memory in accordance with constraints of an operating system 90 and the usage of the main memory 16 by the DMA controller. For example, in the present instance, the command buffer is defined as an 8K-byte page arranged on an 8K-byte boundary. Figure 3 also illustrates the base, tail and head registers 76, 78 and 80 which are held in the DMAC 36. It will be noted that the processor maintains a mirror base pointer 105, a mirror head pointer 109 and a mirror tail pointer 107 in a mirror base register 106, a

mirror head register 110 and a mirror tail register 108, respectively.

[0038] In an embodiment of the invention, the command buffer base register (base register) 76 in the mover 64 of the DMAC 38 enables addressing of the base 91 of the command buffer 92 in main memory 16. Figure 4 illustrates the base register 76. This is a 30-bit register with bits 0-12 being 0 (i.e. at the 8K boundary) and bits 13-29 addressing the individual page within the main memory 16.

[0039] A tail pointer 77 and a head pointer 79 (held in the tail register 78 and the head register 80, respectively, of the DMAC 38) enable the addressing of the tail 98 and the head 100, respectively, of a sequence of DMA commands held in the command buffer 92. The head of the sequence of DMA commands relates to an end of the sequence at which DMA commands are added to (stored in) the command buffer. The tail of the sequence of DMA commands relates to an end of the sequence at which DMA commands are taken from (read from) the command buffer. As will be described in the following, the command buffer is further configured as a ring buffer (i.e. it is addressed by modulo-addressing).

[0040] The base register 76 contains the SBus DVMA address that the DMAC 38 issues in order to find DMA commands in the command buffer 92.

[0041] The tail pointer 77 in the tail register 78 provides an index of the DMA command currently being processed (e.g., a DMA command 102). Each increment in the tail pointer 77 in the tail register 78 represents a 16-byte increment in main memory, which corresponds to one DMA command. A tail pointer value of zero represents the command being at the address pointed to by the base register 76.

[0042] Figure 5 illustrates an example of the tail register 78. It will be noted that bits 0-3 (i.e., the 16 byte boundary) are zero and bits 4-12 provide the index value within the command buffer 92.

[0043] The head register 80 contains the index of the last DMA command (e.g., DMA command 104), which is valid in the command buffer 92 in main memory 16. The addressing provided by the head pointer in the head register 80 is as described for the tail pointer 77 in the tail register 78. As can be seen from Figure 6, which is an illustration of a head register 80, this has the same format as the tail register 78 illustrated in Figure 5.

[0044] The control of the head and tail pointers 79 and 77 is such as to configure the command buffer 92 as a ring buffer (i.e. modulo-addressing is provided), whereby the pointers wrap round from one end to the other of the command buffer 92. In other words, addresses increase from 0 - n (where n corresponds to the maximum index value) and then return to 0. Of course, it will be appreciated that the indexing could equally be in the reverse direction from n to 0 and then back to n.

[0045] In a situation where the DMA controller is

idle and has completed any ongoing DMA for a current command, if the head pointer 79 and the tail pointer 77 in the head and tail registers 80 and 78, respectively, point to the same location in the command buffer 92, then there are no valid DMA commands in the command buffer and the DMA controller knows that no DMA commands need to be transferred at that time.

[0046] The way in which the head and tail registers 80 and 78 are used differs as will be explained in the following. In particular, the head register 80 is managed by the processor 10, whereas the tail register 78 is managed by the DMAC 38.

[0047] There now follows a general description of the operation of the system of Figure 1, and in particular of the operation of the DMAC 38.

[0048] As explained above, the DMAC 38 comprises a loader 62 and a mover 64 and is operable in response to DMA commands, which are provided by the processor(s) 12, 14 to the DMAC 38. The DMA command can specify any start address for DVMA and I/O bus accesses, and any transfer length up to a given maximum, with no restriction on alignment. The mover 64 will handle the transfer as best as it can, using the widest transfers and longest bursts it can. For reads that

do not end at a convenient address, more data than is specified in the DMA command may be read by the mover 64 at either end of the transfer, to use the fastest possible burst transfers. Excess reads will read no more than up to the next appropriate boundary (e.g., a 64-byte boundary on the SBus and an 8-byte boundary on the I/O bus). Writes are handled precisely. Thus, the mover 64 will use whatever combination of small SBus and I/O bus cycles as appropriate to write the expected data in the fastest possible way.

[0049] When required the loader 62 can update a DMA command in main memory with a completion status. If specified in the DMA command, the mover 64 will also generate an SBus interrupt. The DMA command is guaranteed to have been completed and the write to memory to have been requested when the interrupt reaches the processor 10, but the timing relationship between the completion status reaching main memory and the interrupt reaching the processor is not defined.

[0050] During DMA, processors can queue DMA commands in the command buffer 92 in main memory 16 and can signal their presence to the loader 62 with a single processor I/O write, namely the writing of the new head pointer to the head register 80. By analyzing any difference between the head pointer 79 and the tail pointer 77, when the mover 64 completes DMA for one command, the loader can immediately load the next DMA command into mover registers 74 and start DMA afresh, without waiting for any processor support. In other words, the loader 62 uses a difference between the head pointer 79 and the tail pointer 77 to indicate that DMA commands are waiting to be downloaded from the command buffer 92.

[0051] Figures 7A and 7B are flow diagrams illus-

trating the operation of the processor 10 and the DMAC 38.

[0052] Figure 7A represents the steps performed by the processor 10.

[0053] In step S1, the processor writes a new DMA command (e.g. DMA command 103) into the command buffer 92. The processor writes a DMA command to the command buffer 92 at an offset identified by the content of the mirror head register 110 (i.e. the mirror head pointer 109). The base of the command buffer is identified by the content of the mirror base register 106 (i.e. the mirror base pointer 105).

[0054] In step S2, the processor increments the mirror head pointer 109 in the mirror head register 110. The next location within the command buffer 92 is identified by incrementing the content of the mirror head register 110 by 16 bytes so that it points to the next location beyond the currently last DMA command 104.

[0055] In step S3, the processor compares the content of the mirror head register 110 to that of the mirror tail register 108.

[0056] If they are equivalent (in the present embodiment, if they are equal), then this means that there may be no more space in the command buffer. In order to check this, it is necessary to make sure that the content of the mirror tail register 108 corresponds to that of the tail register 78. Accordingly, in step S4, the processor copies the current content of the tail register 78 to the mirror tail register 108. If the DMAC 38 has changed the tail pointer 77 in the tail register 78 since the last time the mirror tail register 108 was updated, then the content of the mirror head register 110 and that of the mirror tail register 108 should no longer be equivalent. The loop formed by steps S3 and S4 is repeated until the test in step S3 indicates that the content of the head register 110 and the tail register 108 is not equivalent.

[0057] When the test in step S3 indicates that the mirror head pointer 109 is not equivalent to the mirror tail pointer 107, then the processor decides in step S5 whether there are any more DMA commands to be written to the command buffer 92. If there are, then control passes back to step S1. If not, then in step S6, the processor copies the content of the mirror head register 110 to the head register 80 in the DMAC 38 for updating the DMAC's head pointer 79 to correspond to the mirror head pointer 109.

[0058] It can thus be seen that the processor 10 may add one or more DMA commands at the head of the buffer 92, and then indicate that all the commands are available to the DMAC 38 with a single write to the DMAC head register 80. This helps to minimize the I/O bandwidth requirements.

[0059] The processor 10 does not need to mirror the DMAC tail pointer 77 at every instant in order to determine whether space exists in the command buffer 92. Using steps S1-S6 described above, the processor 10 can use the mirror tail pointer 107, held in the mirror tail register 108, in most cases. The processor 10 deter-

mines whether there is space in the command buffer 92 by comparing the contents of the mirror head register 110 and mirror tail register 108. If the mirror pointer values in these mirror registers 110 and 108 are not equivalent (e.g., in the present example they are not equal), then there is space in the command buffer 92 at the address pointed to by the mirror head pointer in the mirror head register 110. The processor updates the mirror head pointer 109 itself when it writes to the head register 80 in the DMAC 38, but does not normally modify the mirror tail pointer 107. When a predetermined relationship between mirror head and tail pointers 109 and 107 is detected (in the present instance when there is equality between the mirror head and tail pointer values), the processor 10 checks whether the DMAC 38 has updated the tail pointer 77 since it last checked this. The processor 10 can execute a single read access to the tail register 78 in the DMAC 38 to find the current value of the tail pointer 77 and to copy, or write this to the mirror tail register 108. If the DMAC 38 is handling requests at approximately the rate at which the processor 10 generates DMA requests, the processor 10 will only need to access the DMA tail register 78 once every complete rotation of the ring buffer 92. This further minimizes I/O bandwidth requirements.

[0060] Figure 7B represents the steps performed by the DMAC 38.

[0061] In step S11, the loader 62 is operable to compare the head and tail pointers 79 and 77. If the head and tail pointers 79 and 77 are equivalent, then it is determined that there are currently no DMA commands awaiting transfer from the main memory to the mover 64. Accordingly, the DMAC loader 62 waits until there is a difference between the head and tail pointers 79 and 77. Here it is to be understood that equivalence means that the pointers are effectively pointing to a common DMA command storage location in the command buffer. In the present embodiment, this means that they have the same value. However, in another embodiment they could have values with a predetermined relationship to one another, depending upon the addressing mode(s) used.

[0062] If the head and tail pointers 79 and 77 are not equivalent (i.e., in the present embodiment if they have different values), then, in step S12, the loader 62 uses the tail pointer 77 to access the last command in the command buffer 92 by a DMA read operation. Accordingly, the DMAC loader 62 is able to transfer the next DMA command (e.g. DMA command 102) from the main memory 16 to the mover registers 74.

[0063] Following step S13, in which the mover completes moving data specified in the DMA command, in step S14 the loader is operable to write a "done" or "completion" indication to the command buffer location from which the DMA command has just been read. The processor can use this to monitor which DMA data have been transferred successfully by the DMAC 38.

[0064] In step S15, the DMAC loader 62 updates

the tail pointer 77 to point to the new tail within the command buffer 92.

[0065] Steps S11 to S15 above define functions provided by the loader controller 84 of Figure 1.

[0066] As mentioned above, the head and tail pointers 79 and 77 can be operable to configure the command buffer 92 as a ring buffer by means of modulo-addressing. In this manner, the string of commands awaiting transfer from the main memory to the DMAC 38 can wrap around within the memory space allocated for the command buffer. Effectively, therefore, the modification of the head and tail pointers 79 and 77 is effected in a "modulo"-manner based on the memory space allocated for the command buffer 92.

[0067] As described above, the processor 10 can be operable to write a sequence of one or more DMA commands to the command buffer 92. The sequence of commands can be written, with only the last address of the DMA commands transferred being sent to the head register 80 in the DMAC 38. The tail pointer 77 is updated for each DMA command since it is used to form the address for the next command. The DMAC loader 62 can be operable to transfer a series of DMA commands from the command buffer 92 to the DMAC mover 64. In this case, an interrupt can be generated, or the last command completion status in the main memory can be updated, after the sequence of DMA commands has been transferred from the command buffer 92 to the DMAC mover 64 and the last DMA command has been completed.

[0068] There follows a brief description of the further registers provided in the DMAC loader 62 and the DMAC mover 64.

[0069] The further registers 82 in the DMAC loader 62 can include loader status registers indicating a current loader status. A loader moves between two major states, namely passive and active. A power-on reset, a synchronized reset, any data transfer error, or a direct processor command can move the loader to a passive state. Only a direct processor command can move the loader to the active state.

[0070] In the active state, the loader may or may not be processing DMA commands. When all commands have been completed, the content of the head register 80 (i.e. the head pointer 79) will be equivalent to that of the tail register 78 (i.e. the tail pointer 77). The loader will wait indefinitely in the active state for the processor 10 to write a new value to the head register 80, after which the loader will process new DMA commands. The active state is the normal state for the loader.

[0071] The mover registers 74 in the DMAC mover 64 comprise a number of different registers.

[0072] A DMAC_ACTION register defines the current mover operations. This register is loaded automatically by the DMAC loader 62 and depends on an action byte of a DMA command in main memory. The DMAC_ACTION registers include command and status information. This can include a direction indicator, a

validity indicator, an interrupt enable indicator and a writeback enable indicator. The direction indicator indicates the direction of transfer. The validity indicator enables checking of the validity of the command. The interrupt enable indicator indicates whether the DMA mover generates an interrupt on completion of a DMA operation. The writeback enable indicator indicates whether a completion status is to be written back to the command buffer after completing a command. When the DMAC mover has completed an action command, it updates DMAC_ACTION with a termination status. If write-back is enabled, in the DMAC_ACTION register, the DMA loader writes the updated DMAC_ACTION back into main memory 16. The DMAC_ACTION register may contain further status bits for further status indicators.

[0073] A DMAC_DVMA register defines an SBus DVMA address to be used for DMA. This DVMA address is translated into a real address by an I/O memory management unit in the system I/O controller 30. This register is loaded automatically by the DMAC loader 62 from a DMA command in the command buffer 92.

[0074] A DMAC_ILOAD register defines a I/O bus DVMA address for use for DMA. This register is also loaded automatically by the loader from a DMA command in the command buffer 92.

[0075] A DMAC_COUNT register provides a DMAC transfer byte count. This provides the number of bytes to transfer for a DMA command. This register is also loaded automatically by the DMAC loader 62 from a DMA command in the command buffer 92.

[0076] As mentioned above, the processor 10 maintains mirror base, tail and head pointers 105, 107 and 109 in mirror base, tail and head registers 106, 108 and 110, respectively. As shown in Figure 3, these registers can be held in main memory. Alternatively, they could be formed by CPU registers.

[0077] Figure 8 illustrates further aspects of the DMAC 36 in more detail. In particular, Figure 8 illustrates more details of the DMA transfer mechanism 85.

[0078] As illustrated in Figure 8, a first SBus interface 202 is connected to receive data from the SBus 32. A second SBus interface 204 is connected to transfer data to the SBus 32. The SBus interfaces are clocked by clock signals CS supplied by the mover controller 86. As illustrated in Figure 3, the SBus is a single bus 32. However, it will be appreciated that the SBus may comprise control, address and data lines, and the SBus interfaces 202 and 204 can be configured as multiple clocked buffers, each buffer being responsible for one or more of the control, address and data lines. The output of the SBus interface 202 is supplied via a line 206 to form a first input to a multiplexor 210. A second input to the multiplexor 210 is formed by a path 208 connected to the output of an accumulator 240.

[0079] Selection of the first and second inputs to the multiplexor 210 is effected by multiplex signals MS output by the mover controller 86. The output 211 of the

multiplexor 210 is supplied to a buffer stage 212. The buffer stage 212 includes clocked input gates 214, a buffer storage (e.g. a random access memory with a capacity of 128 by 8 bytes) and clocked output gates 218. The clocked input and output gates 214 and 218 receive clock signals CS from the mover controller 86.

[0080] The output of the buffer stage 212 is supplied to form a first input 226 to a second multiplexor 230, and also an input 224 to a holding register 222. The holding register 222 is also clocked by clock signals CS from the mover controller 86. The output of the holding register 222 forms an input 220 to the second SBus interface 204.

[0081] A second input 228 to the second multiplexor 230 is formed by the output from a first I/O Bus interface 250.

[0082] Selection of the first and second outputs from the second multiplexor 230 is also effected by multiplex signals (MS) from the mover controller 86. The output 232 of the second multiplexor 230 is supplied to a shifter/rotator 234. The shifter/rotator 234 is a controllable barrel shifter for providing a selectable barrel shift. The shifter/rotator 234 receives shift control signals 236 from the mover controller 86. The output 238 from the shifter/rotator 234 is supplied to an accumulator 240. The accumulator 240 receives byte enable signals 242 from the mover controller 86 and also phase signals 244 from the mover controller 86. The output from the accumulator 240 forms the second input 208 to the first multiplexor 210, and also forms an input 246 to a second I/O bus interface 248. The first I/O bus interface 250 and the second I/O bus interface 248 are also clocked by clock signals CS from the mover controller 86. An output 252 from the I/O bus interface 248 is provided for putting data onto the I/O bus 40. An input 254 to the first I/O bus interface 250 is provided for receiving data from the I/O bus 40.

[0083] A description of the operation of the mover transfer mechanism 85 will be described in the following. It will be noted that the mover controller 86 is operable in response to data in the mover register 74 of the mover 64 to control operation of the mover transfer mechanism 85.

[0084] Figure 9 is a schematic representation of the shifter/rotator 234. As illustrated in Figure 4, the input 232 to the shifter/rotator 234 and the output 238 from the shifter/rotator 234 are multi-byte outputs. The shift inputs 236 from the mover controller 86 are operable to control a barrel shift operation within the shifter/rotator 234 selected by the shift control signals 236.

[0085] Figure 10 illustrates the accumulator 240. The inputs 238 from the shifter/rotator 234 to the accumulator 240 are supplied in parallel to a byte mask 239. The byte mask is operable to select one or more bytes from its input in accordance with byte enable signals 242. An accumulator register 241 is operable selectively to accumulate and output the bytes supplied by the byte mask 239 in accordance with phase signals 244. It

should be noted that although in the present instance the masking is on a byte-by-byte basis, in other embodiment, the selection could be on another basis, for example on a bit-by-bit basis. The phase signals control different phases of operation of the accumulator register 241 to accumulate and then output the accumulated bits. The byte enable signals 242 and the phase signals 244 are supplied by the mover controller 86. The outputs from the accumulator register 241 are supplied to the lines 208 and 246.

[0086] In Figure 8, all of the internal data path lines are 64 bits wide (i.e., 8 bytes wide), as marked by a diagonal slash. It should be noted, however, that the data and address lines on the buses 32 and 40 may be controlled separately, by means of separate bus interfaces, in which case the lines connecting the buses to the interfaces may be other than the internal data width. It will also be appreciated that the internal data width in the present example is merely one example of a possible internal data width, and other internal data widths (for example, 32 bits or 128 bits) may be used instead.

[0087] Figure 11 is an illustration of a possible transfer between a source bus (which can be either the SBus 32 or the I/O bus 40) and a destination on the other of those buses. In the present instance, an 8 byte-wide data source is shown as the embodiment provides 64 bit-wide data paths. As suggested above however, in other embodiments other data widths may be employed. Of course the DMA transfer could be between other buses depending on the location of the DMAC 38.

[0088] It will be noted that the source information 260 starts at byte 4 of word 0. The destination information, however, starts at byte 6 of word 10. Accordingly, it can be seen that it is not possible merely to transfer the bytes word-by-word from the source to the destination. The mover transfer mechanism is arranged to provide very efficient transfer of data from the source to the destination addresses.

[0089] There follows a description of the operation of the shifter/rotator 234 and the accumulator 240, with reference to Figure 12.

[0090] In Figure 12, the first column represents the inputs 232 to the shifter/rotator 234 for first, second, third and fourth stages in the transfer. The second column indicates the output 238 of the shifter/rotator 234 in those four stages. The third column represents the byte enable signals 242 supplied to the accumulator 240 in first and second phases for the stages 1-4. The fourth column illustrates the content of the accumulator register 241 in each of the phases of the four stages. The arrows in the fifth column represent the output at 245 of the content of the accumulator 240 during the first phase of each of the four stages 1-4. The references in the first column 1.1, 1.2, 2.1 etc., represent the first phase of the first stage, the second phase of the first stage, the first phase of the second stage, etc. A more detailed explanation of Figure 12 will now follow.

[0091] Accordingly, in the first stage (as shown at 1.1) the input to the shifter rotator 234 is the content of the source word 0 (i.e. ***ABCDE) The shifter/rotator 234 barrel shifts this information to the right by 2 byte positions to give the signals indicated in column 1.1 at 238 (i.e. DE***ABC). This information is immediately available at the output of the shifter/rotator 234. At the same time, in this first phase, the byte enable signals illustrated in row 1.1 at 242 (00000111) are supplied to the byte selector 239 of the accumulator 240. This accordingly causes the accumulator 240 to accumulate in the accumulator register 241 the bytes available at the positions identified by a "1" in the byte enable signals, and not to select the content of the byte positions corresponding to a "0" in the byte enable signals. The content of the accumulator register 241 is immediately output (as illustrated in row 1.1 at 245) to form the first destination byte 10 (i.e. *****ABC).

[0092] Three bytes are shown as being output in one go. In other embodiments, where this is not desired, or possible, the mover controller 86 could supply signals to cause the three bytes to selected one at a time for output.

[0093] In a second phase of operation as illustrated in row 1.2, the inputs and outputs to and from the shifter/rotator 234 remain the same. However, the byte enable signals 242 supplied to the accumulator change to 11000000. As a result, this causes the first two bytes (D and E) input to the accumulator 240 to be selected by the selector 239 and to be accumulated in the accumulator register 241. However, these signals are not output, but are instead held in the accumulator register until the next phase of operation.

[0094] The next phase of operation forms the first phase of step 2. As shown at row 2.1, the input 232 to the shifter/rotator 234 changes to the information for word 1 from the source bus (i.e., FGHIJKLM). The output 238 from the shifter/rotator 234 corresponds to the input barrel shifted to the right by two bytes (i.e., LMF-GHIJK). In this phase, the byte enable signals are changed to 00111111 as shown at 242 in Figure 7, whereby the selector 239 selects the signals FGHIJK for accumulation in the accumulator unit 241. This then completes the second word to be supplied to the destination, and this is then output as a word of 8 bytes as indicated by the arrow at 245 in column 2.1.

[0095] It will be appreciated from study of Figure 12 how this process continues for phases 2.2, 3.1, 3.2 and 4.1. It will be noted that there is no need for a phase 4.2 as the information to be output for the fourth destination byte (byte 13) is already contained in the accumulator at the end of phase 4.1. Accordingly, in this example, this information is immediately output as shown by the arrow at 245 in column 4.1. In other examples, as less than a word of eight bytes is to be output, this could be output in groups of one, two and/or four bytes according to a particular implementation.

[0096] It can, accordingly, be seen how information

can very effectively be barrel shifted and accumulated to transfer from a source alignment to a destination alignment for transfer between the buses 32 and 40, or between the buses 40 and 42.

[0097] The detailed movements of bytes can vary from implementation to implementation. For example, when both buses are operating in burst mode, the accumulator 240 can be caused to operate in two phases as described above. The first phase captures n bytes of an I/O data word x from the shifter. The second phase captures (8-n) bytes from an I/O data word (x+1). The content of the accumulator is then transferred to the buffer 216 and immediately captures n bytes from (x+1) and so on. The value of n is related to the shift and depends on the byte offsets of the start address on the source and destination buses.

[0098] It is noted that the shift signals 236, the byte enable signals 242 and the phase signals 244 are all provided by the mover controller 86. The mover controller 86 is operable to respond to data stored in the mover registers 74, including the source and destination addresses, and the direction of transfer, to generate these signals, and also the clock CS and multiplex signals MS supplied to the other units illustrated in Figure 8.

[0099] Figure 13 is a block diagram summarizing the operation of the mover controller. Accordingly, in step S21 the mover controller retrieves source and destination addresses from the mover register 74. In S22, it determines the direction of transfer between the buses 32 and 40. This information is typically determined by a transfer direction indication in one of the mover registers 74. However, it could alternatively be derived directly from the source and destination addresses where possible using the bus protocols concerned.

[0100] In S23, the mover controller determines the shifting to be performed within the shifter/rotator, and a sequence of block sizes for transfer.

[0101] It should be noted that this information can readily be calculated in accordance with well established algorithms, given the exact source and destination addresses and the direction of transfer. The determination of the blocks for transfer is established to effect as efficient as possible a transfer between the buses.

[0102] As illustrated in Figure 14, the transfer can be effected for different block sizes depending on the availability for block transfers of those sizes to be effected within the available bus protocols and the available hardware design width. This can be freely chosen.

[0103] As illustrated in Figure 14, it is envisaged that the transfer blocks may have a single byte size, an 8-byte size and a 64-byte size. In other embodiments, bit transfers would also be possible for amounts less than one byte. As represented schematically in Figure 14, a theoretical transfer could include:

[0104] In step S31 initial transfers with byte size blocks is effected until an "8-byte" block boundary is reached.

[0105] In step S32, bytes could be transferred on an 8-byte block size until a "64-byte" block boundary is reached.

[0106] In step S33, transfers could be effected on 64-byte blocks until, approaching the end of the transfer, it is necessary to return to 8-byte block sizes.

[0107] In Step S34, 8-byte blocks are transferred until it is necessary to transfer the remaining bytes as byte sized blocks.

[0108] In step S35, the remaining bytes are transferred.

[0109] As indicated above, a first step prior to step S31 and/or a final step after step S35 could provide transfers on a bit basis.

[0110] It can be seen from the above that the determination of a combination of one or more of the stages above can readily be effected by the mover controller 86, given the size of the transfer to be effected and the exact source and destination addresses, as distributed with respect to 8-byte and 64-byte block boundaries. Clearly, if a suitable alignment is already effected, one or more of the steps illustrated in Figure 9 can be omitted.

[0111] Returning to Figure 13, once the determinations of Step S23 have been made, the detailed sequence of control signals can be generated in step S24 in order, in step S25, to execute the transfer.

[0112] As mentioned above, it is necessary to determine the direction of transfer between the buses 32 and 40. Thus, the transfer may be from the SBUS 32 to the I/O bus 40, or vice versa.

[0113] Control of the direction of transfer is effected by providing the multiplex signals MS to the first multiplexor 210 and the second multiplexor 230. Figure 15 illustrates the transfer of data from the I/O bus 40 to the SBUS 32. In order to effect this, the first multiplexor 210 is controlled to select its second input, that is the output of the accumulator 240. The second multiplexor 230 is controlled to select its second input 228 from the I/O bus interface 250.

[0114] Accordingly, bytes to be transferred received at the I/O bus interface 250 from the I/O bus 40 are supplied to the second multiplexor 230, which in turn supplies those bytes to the shifter/rotator 234. The shift/rotation and accumulation process already described above with respect to Figures 6 and 7 is effected in the shifter/rotator and accumulator and the bytes output from the accumulator are supplied to the first multiplexor 210. The output from the first multiplexor 210 is buffered in the buffer 216 to form blocks of data for transfer to the holding register 222 for transfer via the SBUS interface 240 to the SBUS 32. The buffer 212 and the holding register 222 enable buffering of data between the I/O bus and the SBUS 32 to take account of different bus speeds and bus access availability.

[0115] It will be noted that the lines 210, 206, 226, 246 and 252 are shown dashed in Figure 15. This is

used to represent that data is not being transferred along those paths in a situation where data is transferred from the I/O bus 40 to the SBUS 32.

[0116] Figure 16 illustrates the alternative situation that data transfer is being effected between the SBUS 32 and I/O bus 40. In this case, it will be noted that the lines 201, 206, 226, 246 and 252 are shown as solid lines, whereas the lines 254, 228, 208, 224, 220 and 203 are shown dashed. This is to indicate that data does not flow along the dashed lines during the transfer of information from the SBUS 32 to the I/O bus 40.

[0117] Accordingly, in the situation illustrated in Figure 16, data is received from the SBUS 32 over the lines 201 to the SBUS interface 202. The data from the SBUS interface 202 is supplied via lines 206 to the first multiplexor 210. In this case, multiplex signals MS are supplied to the first multiplexor 210 to select its first input. Accordingly, the data received at the input 206 of the multiplexor 210 is supplied to its output 211 for buffering in the buffer unit 212. The data is output from the buffer unit 212 in order to provide a constant flow through the shifter/rotator and accumulator 240 for transfer to the I/O bus 40. Thus, data output from the buffer unit 212 is supplied to the first input 226 of the second multiplexor 230.

[0118] Multiplex signals MS are supplied to the second multiplexor 230 to select its first input 226. Accordingly, the data of the first input 226 of the multiplexor 230 are supplied to its outputs 232. The shift/rotation and accumulation operations described above with respect to Figures 6 and 7 are effected by the shifter/rotator 234 and the accumulator 240. The data output therefrom is supplied via lines 245 and 246 to the I/O bus interface 248 for output via lines 252 to the I/O bus 40.

[0119] Accordingly, there has been described a shift-accumulate mechanism '234:240' which is used in accordance with a FIFO buffer 216 to take data from a source bus and to realign it to a desired alignment for output to a destination bus. In this particular example, the two buses are an SBUS 32 and an I/O bus 40. The source could be either one of those buses and destination the other of those buses. However, it will be appreciated that the present invention is generally applicable to buses operable under any suitable bus protocol.

[0120] In the present instance, the natural byte alignment used on an SBUS is 64 bytes and the mechanism enables the use of 64 byte burst transfers wherever possible. The I/O bus in the present example is 8 byte aligned. In the present example the I/O bus transfers may be up to 512 bytes in length, starting and ending on an 8-byte boundary. However, in other examples different sizes of buses and bus transfers may be effected. In a first mode (DMA read mode), the I/O bus is the source and the SBUS is the destination. The shifter/rotator 234 is an 8 byte barrel shifter which rotates the data depending on the relative byte alignments of the starting source and destination addresses. I/O bus transfers can be used to read from the source address. It will be noted with reference to Figure 12 that

some extra data may be read at the beginning and end of the transfer. Bytes are captured in the accumulator until the destination address is 8 byte aligned. The accumulator contents can then be transferred to the buffer 216. The buffer content can then be written to the SBus using byte, or 8 byte transfers until a 64-byte address boundary is reached, at which point 64 byte burst transfers are used.

[0120] When both buses are in burst mode, the accumulator operates in two phases. The first phase captures n bytes of an I/O data word x from the shifter. The second phase captures (8-n) bytes from an I/O data word (x+1). The content of the accumulator is then transferred to the FIFO buffer and immediately captures n bytes from (x+1) and so on. The value of n is related to the shift and depends on the byte offsets of the start address on the source and destination buses.

[0121] Operation in DMA write mode is similar except that the SBus is then the source and the I/O bus is the destination. SBus reads always begin on 64 byte bursts. Some extra bytes may be read at the beginning and end of the transfer. Data is read into the FIFO buffer 216 and proceeds through the shifter/rotator 234 and accumulator 240 to be written to the I/O bus. Individual bytes (or groups of bytes) are written to the I/O bus until an 8 byte address boundary is reached, at which point 8 byte burst transfers are used until less than 8 bytes remain. Remaining bytes are then transferred singly.

[0122] It can be seen that a very effective and flexible method has been described for transferring bytes between two buses. Although this has been described with reference to transfers on byte, 8 byte and 64 byte boundaries, it will be appreciated that the transfers could be effected on other boundaries, and could be effected on a sub-unit of a byte, for example, a bit basis.

[0123] Although a particular exemplary configuration of a computer system including the invention has been described above, it will be appreciated that the present invention is applicable not only a computer system having such a configuration.

[0124] Indeed, it will be appreciated that although particular embodiments of the invention have been described, many modifications/additions and/or substitutions may be made within the scope of the present invention.

Claims

1. A bridge for a computer system, which computer system includes a first bus and a second bus, the bridge being locatable between the first and second buses and comprising:

shifter means for shifting bytes from one of the first and second buses; and accumulator means for selective accumulation of bytes shifted by the shifting means to form realigned bytes to be passed to the other of the first and

second buses.

2. A bridge according to claim 1, wherein the shifter means has bytes comprises an input connected to receive bytes from one of the first and second buses and an output providing shifted bytes with a selectable shift with respect to the received bytes; and

the accumulator means has an input connected to receive the output of the shifter means and provides selective accumulation of the shifted bytes, the accumulator having an output for supplying realigned bytes to be passed to the other of the first and second buses.

3. The bridge of claim 1 or claim 2, comprising control logic having a first control output connected to the shifter means for control signals to select the shift applied to the received bytes and a second control output connected to the accumulator means for control signals to select part of the shifted bytes to be accumulated.

4. The bridge of claim 3, wherein the control logic has a further output for supplying phase signals during successive cycles of operation for defining a first phase and a second phase during an operating cycle, the shifter means being operable to receive an input set of bytes during an operating cycle.

5. The bridge of claim 3, wherein the accumulator means comprises an output register, the accumulator means being arranged:

in the first phase, to complete accumulation of an output set of realigned bytes by registering the output of selected positions in the shifted input set of bytes at corresponding positions in the output register; and

in the second phase, to start accumulation of a subsequent output set of bytes by registering the output from remaining positions in the shifted input set of bytes at corresponding positions in the output register.

6. The bridge of claim 5, wherein a set of bytes forms a word.

7. The bridge of claim 5, wherein a set of bytes forms part of a word.

8. The bridge of any of claims 5 to 7 wherein the control logic is arranged to determine the control signals for successive cycles and phases to effect a burst transfer of bytes from one bus to the other bus.

9. The bridge of any preceding claim, further comprising a buffer for buffering sets of bytes.

10. The bridge of claim 9, wherein the buffer is connectable between the first bus and the shifter means when bytes are to be transferred from the first bus to the second bus. 5

11. The bridge of claim 9 or claim 10, wherein the buffer is connectable between the accumulator means and the first bus when bytes are to be transferred from the second bus to the first bus. 10

12. The bridge of any of claims 9 to 11, comprising a first multiplexor having a first input connected to receive bytes from the first bus, a second input connected to receive bytes from the accumulator means and an output connected to supply bytes to the buffer. 15

13. The bridge of any of claims 9 to 12, comprising a second multiplexor having a first input connected to receive bytes from the second bus, a second input connected to receive bytes from an output of the buffer and an output connected to supply bytes to the shifter means. 20

14. The bridge of any preceding claim, wherein the bridge forms a direct memory access controller. 25

15. A bridge according to any preceding claim, wherein the bridge is configured to be operable as a direct memory access controller.

16. A computer system having a first bus, a second bus and a bridge according to any preceding claim. 30

17. A method of transferring data between a first bus and a second bus in a computer system, the method comprising:
shifting bytes received from one of the first and second buses to providing a selectable shift to the received bytes; and
performing selective accumulation of the shifted bytes for supplying realigned bytes to be passed to the other of the first and second buses. 35

18. The method of claim 17, operable in successive cycles, wherein, in each cycle, a set of bytes is received from said one of the first and second buses. 40

19. The method of claim 18, comprising:
in a first phase of a cycle, completing accumulation of an output set of realigned bytes by reg- 45

istering a selected part of a shifted input set of bytes at a corresponding part of an output register;
outputting the output set of realigned bytes from the output register; and
in a second phase of a cycle, starting accumulation of a subsequent output set of realigned bytes by registering a remaining part of the shifted input set of bytes at a corresponding part of the output register. 50

20. The method of claim 17, wherein, for transfers from the first bus to the second bus, bytes are buffered prior to shift and accumulation steps, and, for transfers from the second bus to the first bus, bytes are buffered after shift and accumulation steps. 55

21. The method of claim 17 for transferring a burst of bytes between the first and second buses, wherein:
in an initial stage, at least a selected parts of shifted bytes from up to two cycles are accumulated to form an initial set of output bytes;
in an intermediate stage, at least a selected part of shifted bytes from two cycles are accumulated to form intermediate sets of output bytes; and
in a final stage, at least a selected part of shifted bytes from up to two cycles are accumulated to form a final set of output bytes. 60

22. The method of claim 17 for transferring a burst of bytes between the first and second buses, wherein, to achieve realignment with efficient block transfer, a block transfer includes at least one of the following stages:
an initial transfer of selected bits;
an initial transfer of selected bytes;
an initial transfer of selected blocks of bytes;
an initial transfer of sub-blocks of words;
a block transfer of words;
a final transfer of sub-blocks blocks of words;
a final transfer of selected blocks of bytes;
a final transfer of selected bytes; and
a final transfer of selected bits. 65

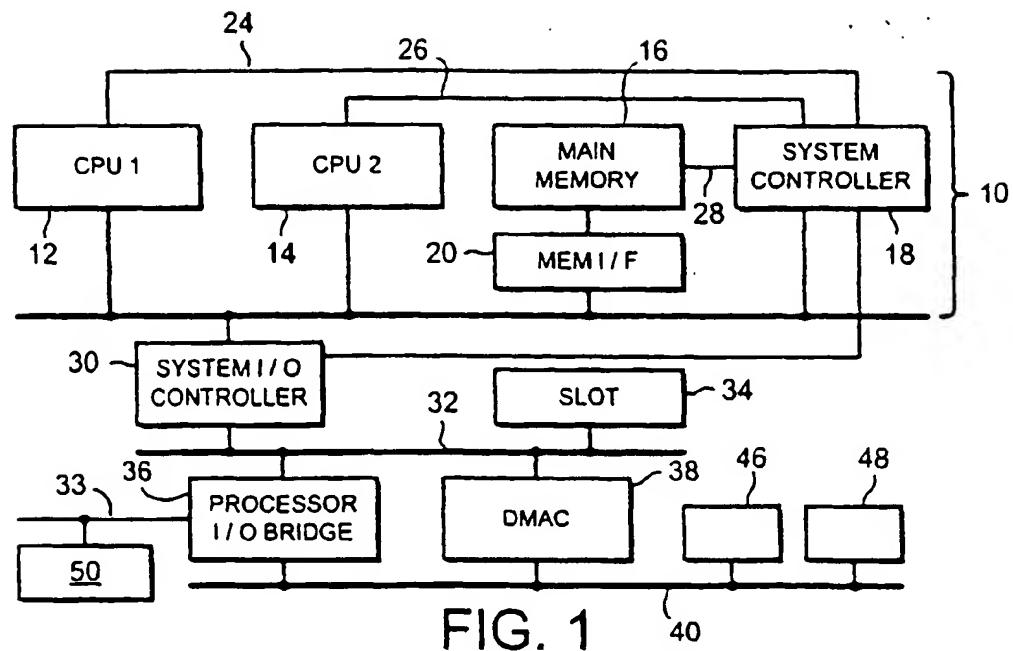


FIG. 1

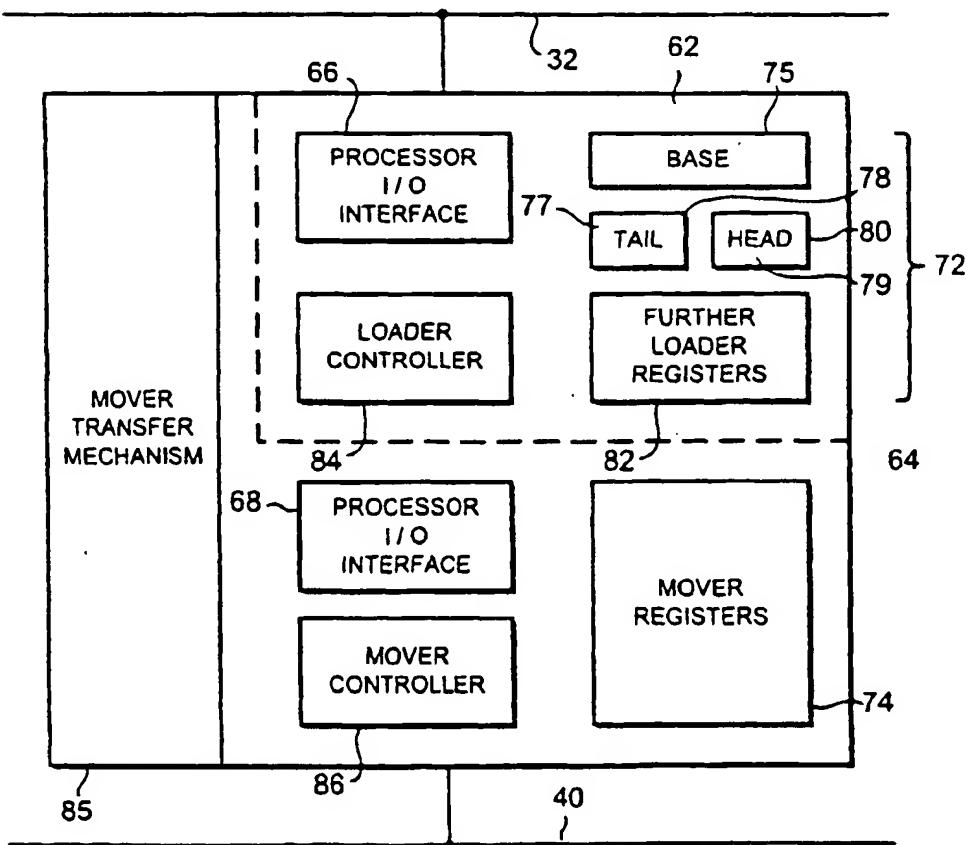


FIG. 2

BEST AVAILABLE COPY

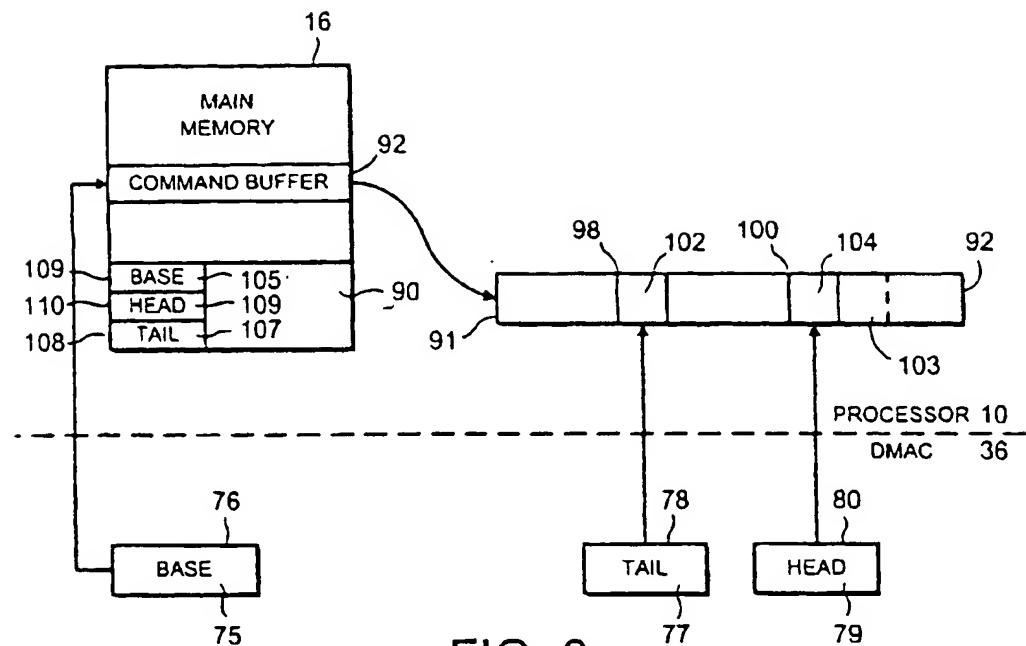


FIG. 3



FIG. 4

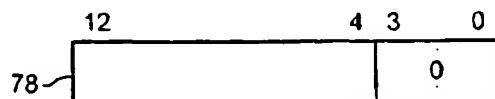


FIG. 5



FIG. 6

BEST AVAILABLE COPY

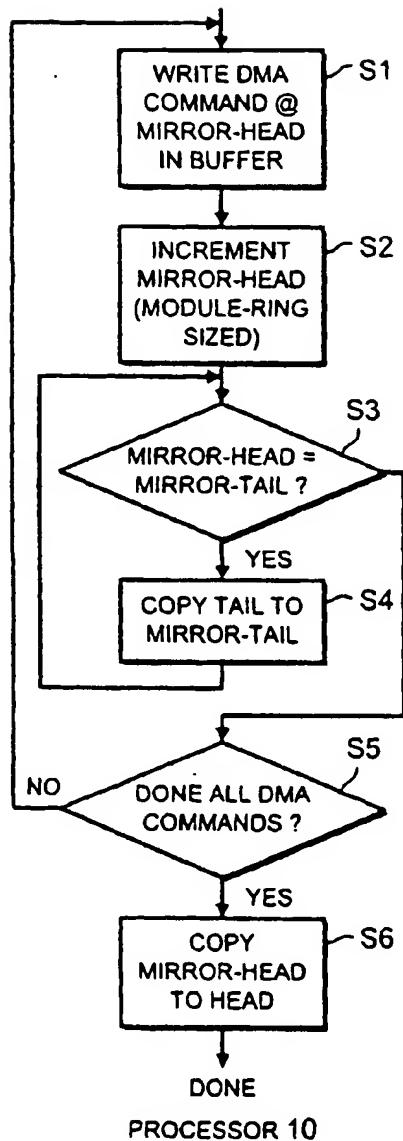


FIG. 7A

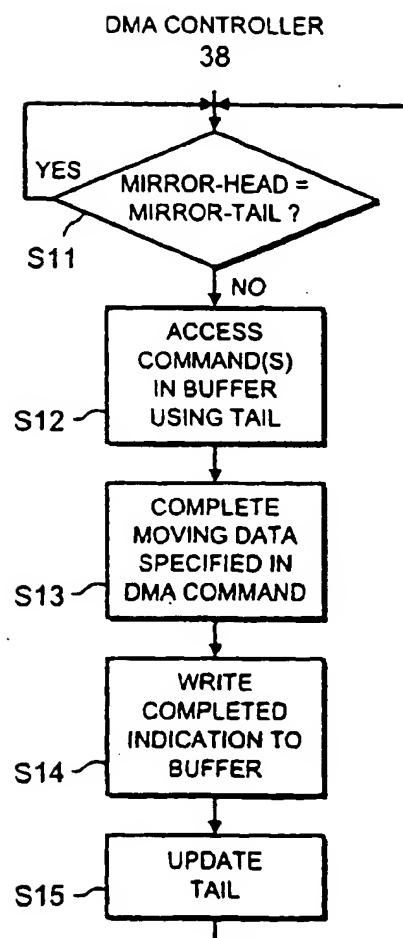


FIG. 7B

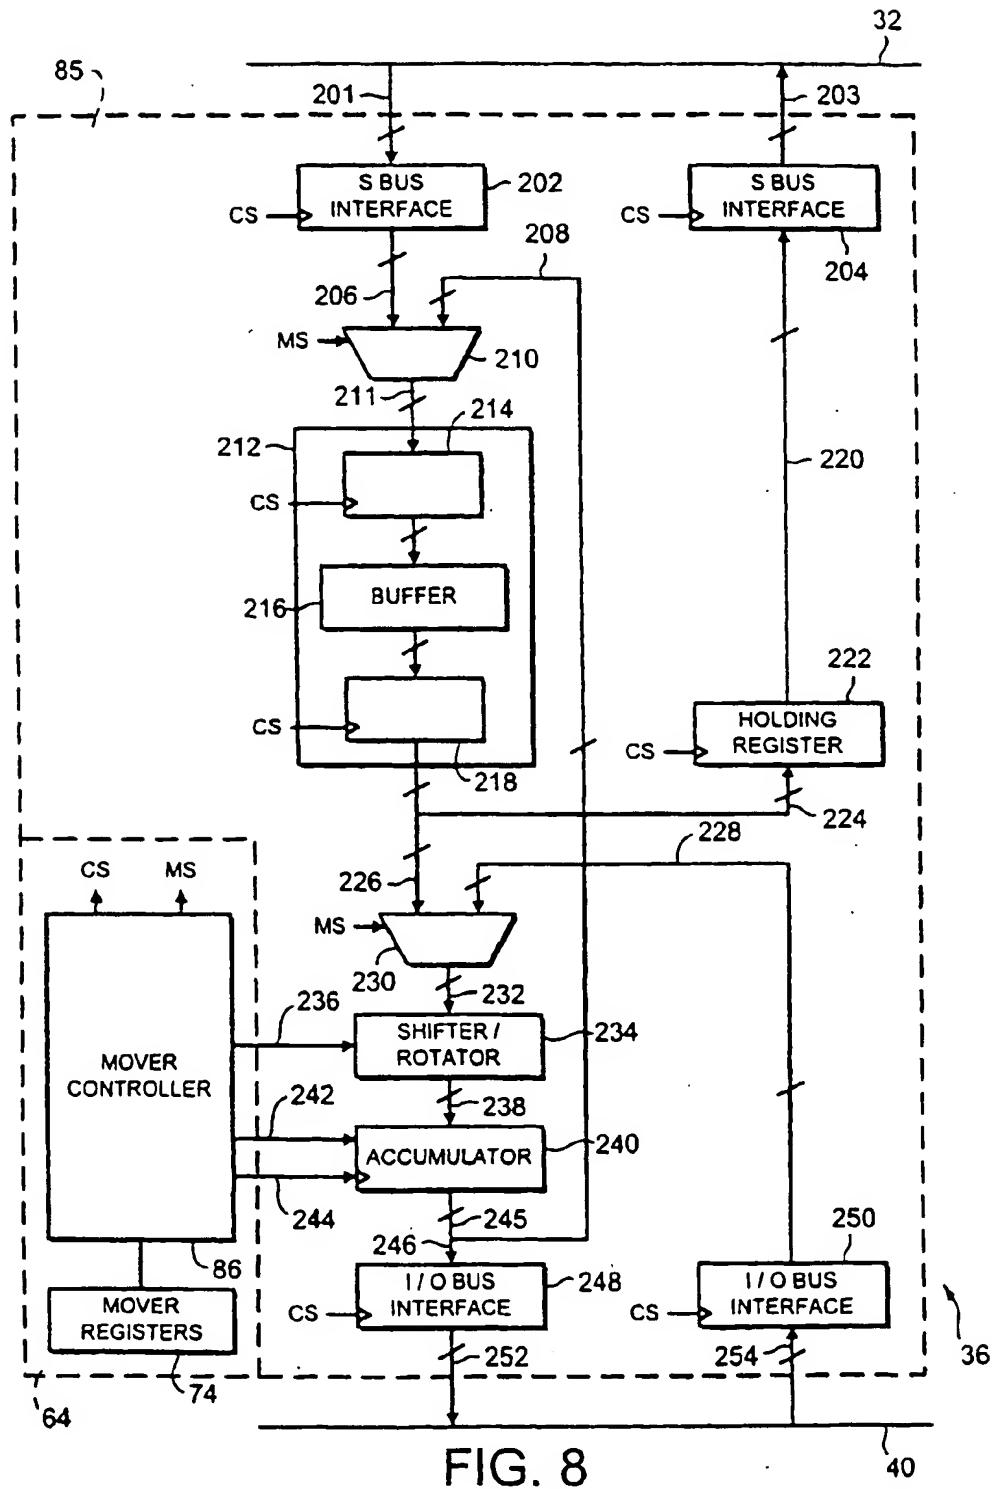


FIG. 8

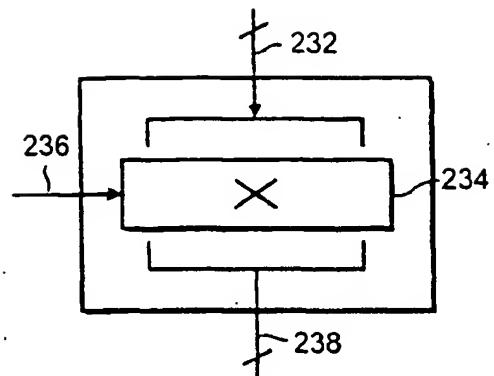


FIG. 9

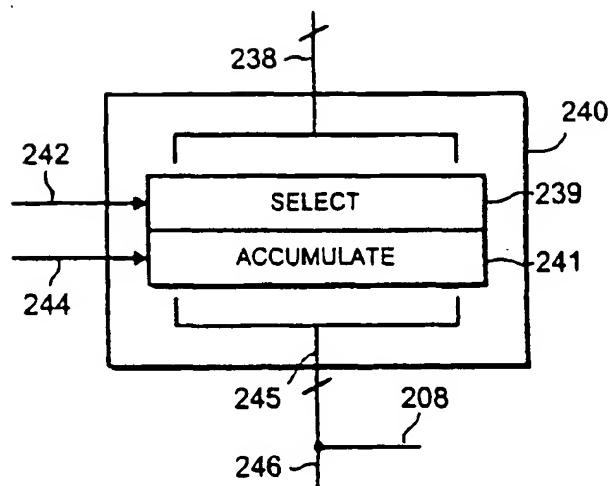


FIG. 10

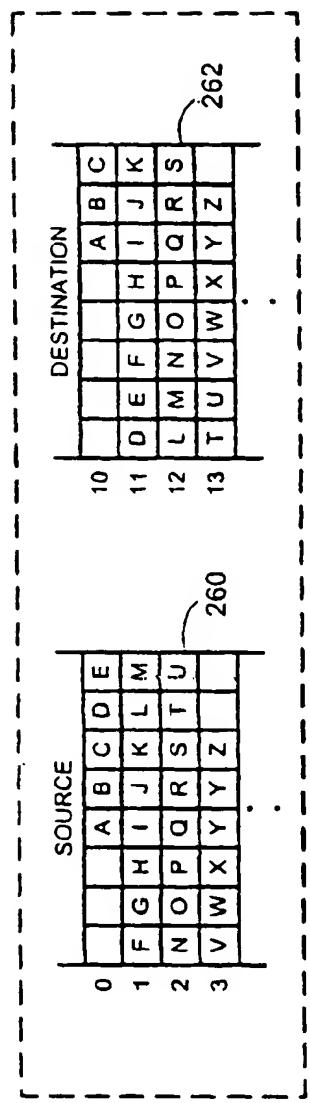


FIG. 11

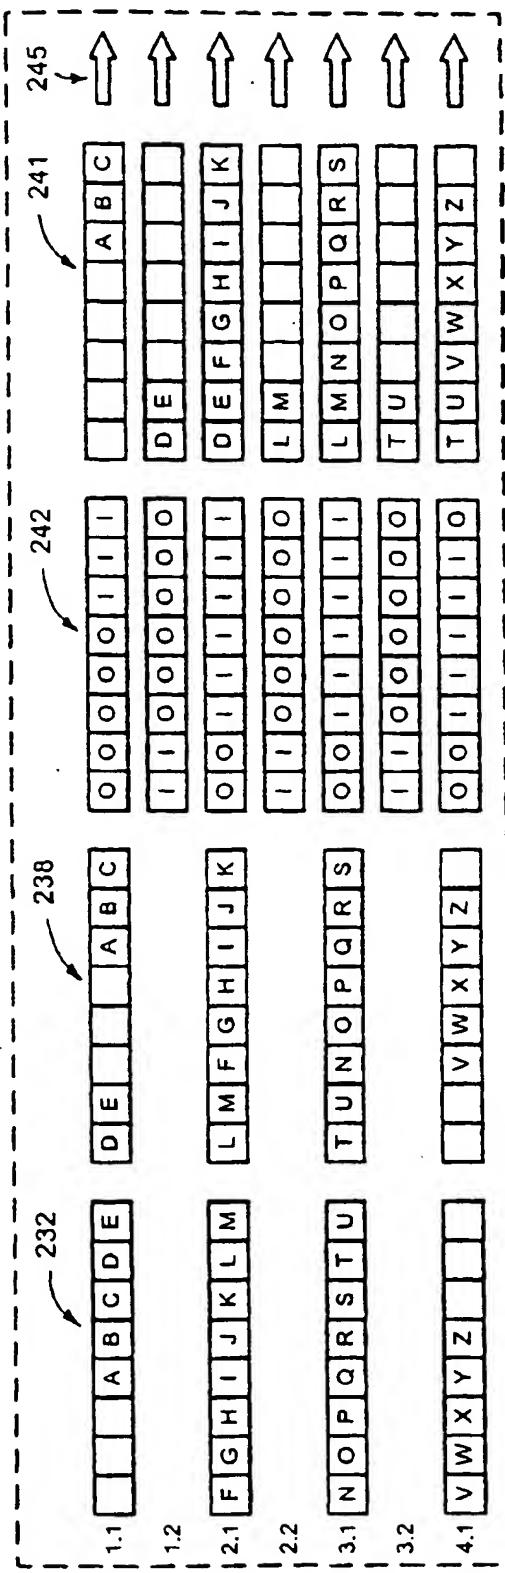


FIG. 12

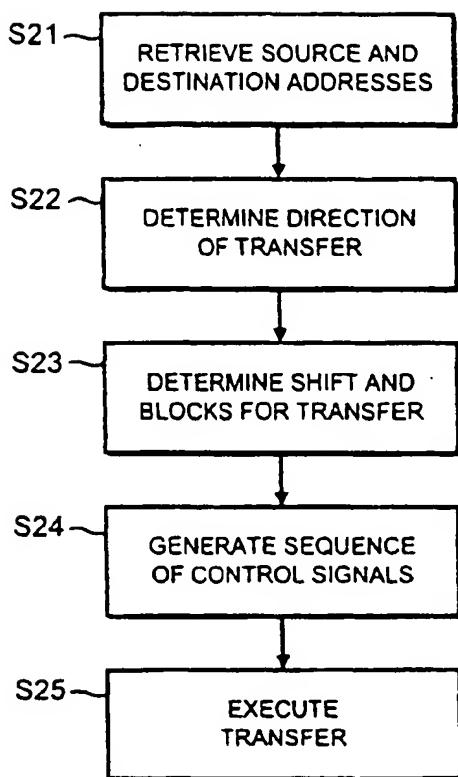


FIG. 13

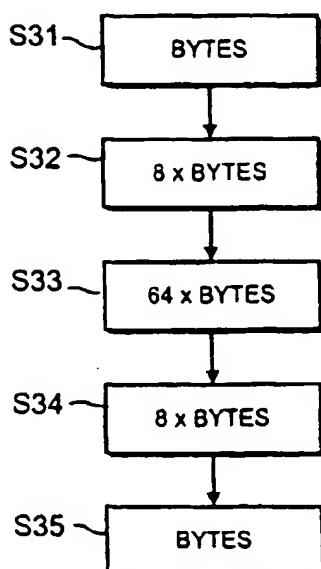


FIG. 14

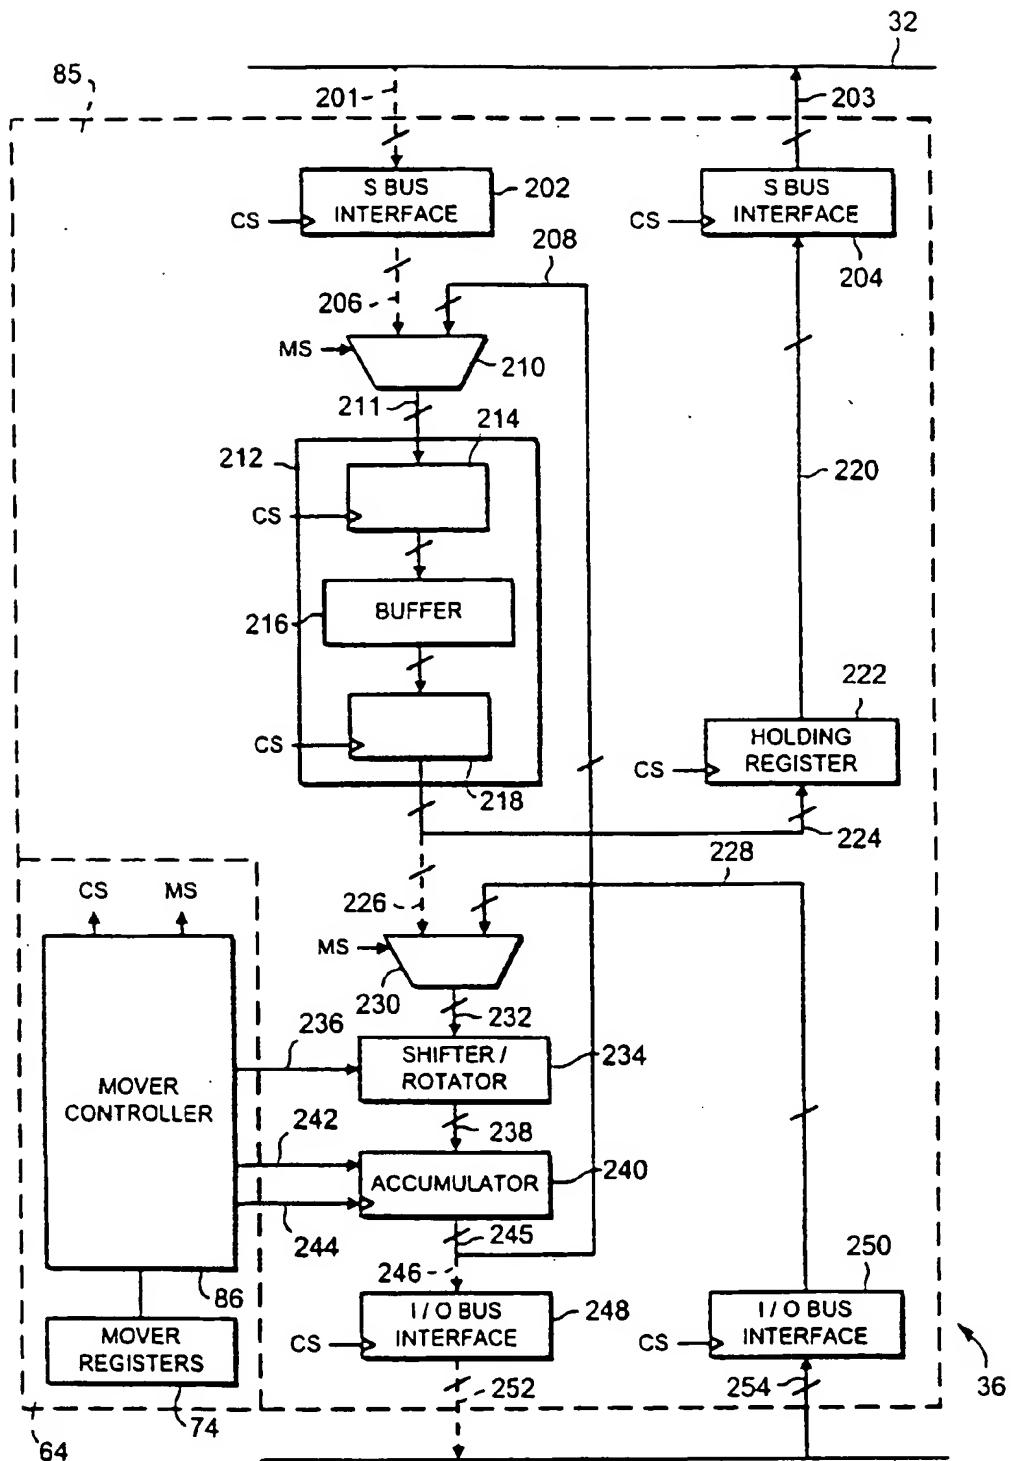


FIG. 15

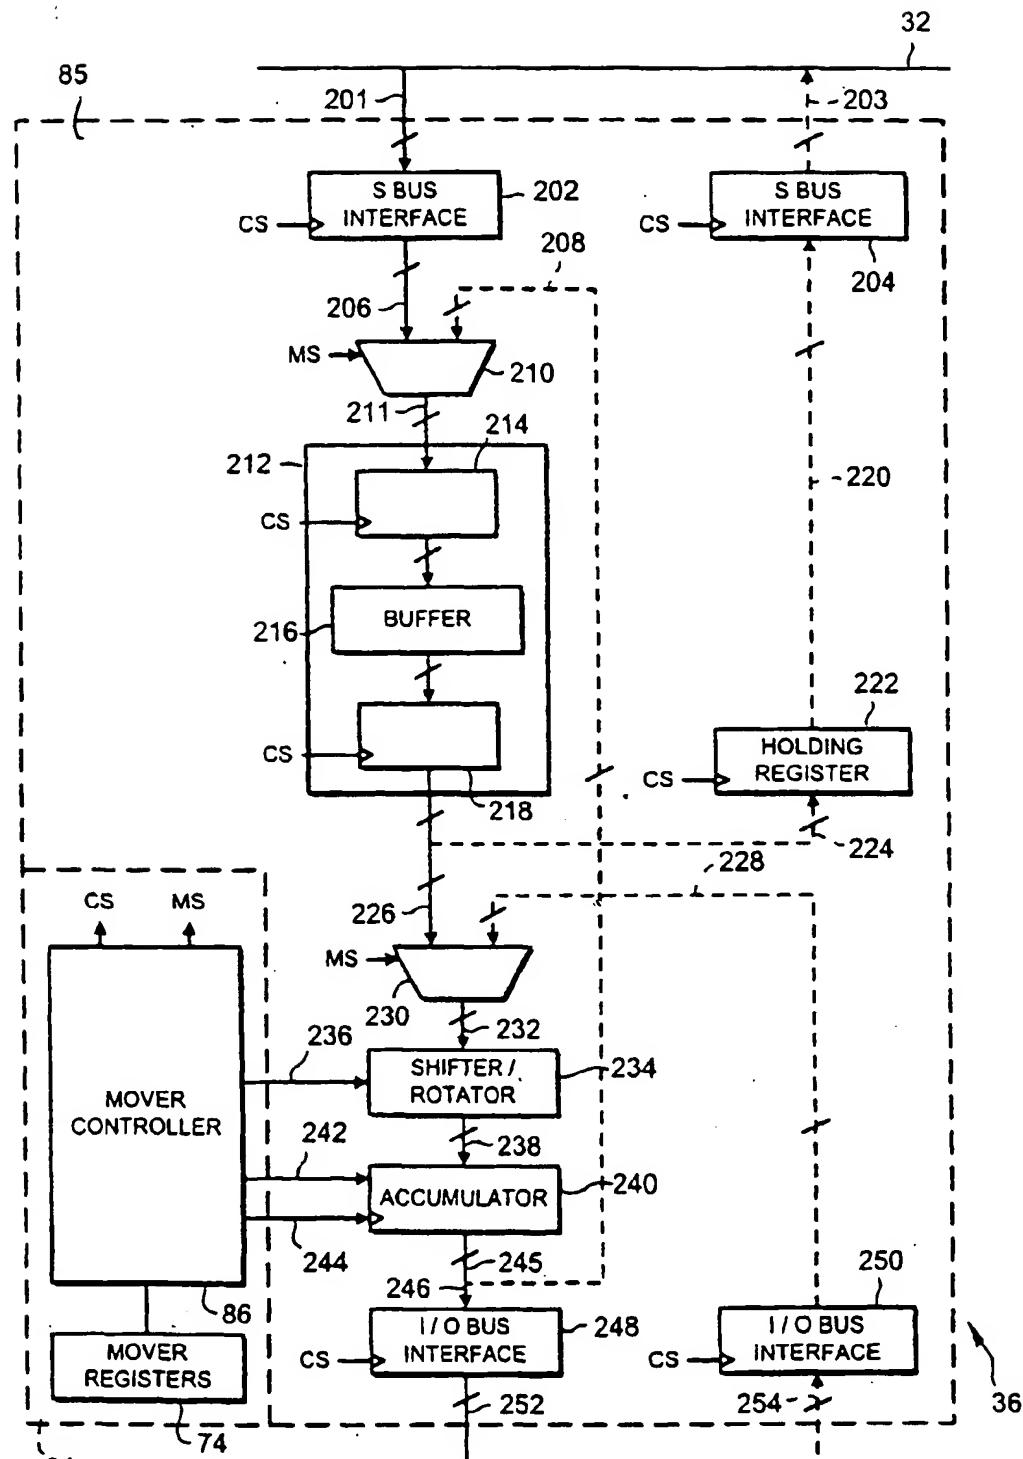


FIG. 16